Computation is essentially medium-independent

Zoe Drayson, UC Davis

Draft, February 2024 – not for quotation

1. Medium independence

There is a widespread idea that computation is essentially *independent* in some important sense from its physical implementation or realization. Several terms have been used to label this feature: computation has been described as 'medium-independent, 'multiply realizable', 'substrate neutral', and 'organizationally invariant'. While these terms can be specified in distinct ways, the differences need not concern us here – I'll use the term 'medium independence' throughout to capture a somewhat general idea. My goal will be to defend the essential medium-independence of computation from recent attempts to argue against it.

How should we understand what Milkowski (2016) refers to as "the intuitive idea that computation is not really linked to the physical substrate in the way many other properties are"?

> "Formal systems are independent of the medium in which they are 'embodied'. In other words, essentially the same formal system can be materialized in any number of different media, with no formally significant difference whatsoever. This is an important feature of formal systems in general. I call it medium independence." (Haugeland 1985)
>
> "To a first approximation, a property or process is medium independent if it can be realized in different physical media." (Curtis-Trudel 2021)
>
> "This view— that representations and algorithms are independent of (or can be analyzed without reference to) their implementing physical media—has been referred to as "medium independence" in the philosophy of computation." (Maley 2021)

Why think that computation is essentially medium-independent? The medium-independence of computation seems to be built in to the very concept of computation associated with computability theory. Computability theory is a branch of mathematical logic which studies computable functions, where a function is computable if it can be specified algorithmically. Any function which can be algorithmically computed can be described in terms of a Turing Machine, which is an idealized machine which can be specified mathematically: its instructions (its machine table) can be represented as a sequence of numbers. Turing Machines themselves are abstractions which do not rely

on concrete computation. This gives us a very clear sense of how computation can be understood as medium-independent, which very few people would deny.

Our everyday digital computers, however, are not themselves abstract objects: they are (approximate) *physical implementations* of an idealized Universal Turing Machine. Similarly, when cognitive scientists claim that the brain is a computer, they are claiming that neural states and processes *physically implement* mathematically-defined computations. This raises an important question about the implementation relationship between abstract computations and physical systems: what is it for a physical system to implement a computation? A number of answers have been suggested. According to mapping accounts, a physical system performs a computation just in case the transitions between the states of the physical system mirror the transitions between the states of the computational function. To avoid triviality, acceptable mappings are restricted by appealing to certain causal properties, dispositional properties, or other counterfactual-supporting properties. These mapping accounts can be further restricted by adding a semantic constraint on which computation also involves the processing of representations. I won't be defending any particular account of computational implementation here: I will be assuming that for a physical process to implement a computation is to for that physical process to be related in some appropriate way to a mathematically-defined abstract computation, but I will remain neutral on the nature of the appropriate relation.[1]

Those who deny the essential-medium independence of computation usually have physical computations in mind, and often some particular subset of physical computations: computation in a specific kind of physical system (e.g. the brain) or a specific kind of physical computation (e.g. non-digital). Here are some representative quotations from those who deny the medium-independence of some subset of physical computation:

> "We argue that the sense in which brains compute does not imply that brains implement multiply realizable computational processes. [...] even if computational explanations and models in the cognitive and brain sciences are to be taken as literally postulating that cognitive processes are computational processes that are realized by brain processes, that does not imply that cognitive or brain processes are multiply realizable." (Polger and Shapiro, 2023)
>
> "most of the information processing happening in the nervous system [...] cannot be divorced from the actual stuff it is made of. [...] neural computation is functional manipulation of a medium-dependent vehicle." (Chirimuuta 2022)

[1] I will be rejecting any purported account of the implementation relation in which mathematical computation plays no role. Following Williams (2023), I agree that an account which attempts to define physical computation entirely in terms of concrete properties of physical systems cannot capture the relation between the physical system and the mathematical computation it implements.

> "For digital representation and computation, this is a well-known and often repeated point: this type of representation and computation is medium-independent in a deep and important sense. However, the extent to which this independence applies to analog representation and computation is unclear; I will argue that it simply does not." (Maley 2021)

In what follows, I will consider three sets of recent arguments which purport to show that computation is not essentially medium-independent.

(i) The first set of arguments (call them *Concreta arguments*) focuses on the concreteness of physical computational systems. Proponents of these arguments generally (a) allow that there is a purely abstract notion of computation on which it is medium-independent, (b) claim that the notion of concrete computation which plays a role in scientific explanation is distinct from the abstract notion, and (c) claim that concete computation is not essentially medium-independent. Against Concreta arguments, I will argue that there is no notion of physical computation which can be specified entirely in concrete terms with no reference to abstraction.

(ii) The second set of arguments (call them *Empirical arguments*) focuses on supposed actual cases of medium-dependent computation in computer science and cognitive science (especially neuroscience). Proponents of these argument argue that (a) we know of actual physical processes which physically implement computations and (b) these actual physical processes are importantly dependent on particular physical properties which another physical medium could lack. Against Empirical arguments, I will argue that the question of whether the physical implementation basis of a computation is specified in medium-specific properties is irrelevant to the question of whether physical computation itself is medium independent.

(iii) The third set of arguments (call them *Analogicity arguments*) focuses on the nature of analog representation. Proponents of these arguments (a) claim there is a kind of representation, analog representation, which is medium-dependent, (b) claim that there is a particular kind of computation, analog computation, which takes place over analog representations, and (c) claim that analog computation is not essentially medium-independent, in virtue of the medium-dependent nature of analog representation. Against Analogicity arguments, I will argue that even if analog representations are medium-dependent, there are fundamental problems with relying on such representations to establish the existence of processes which are both computational and medium-dependent.

2. Concreta arguments

Proponents of Concreta arguments propose that the notion of medium-independence associated with computation applies only to mathematical concepts of computation. They argue that if we are interested in physical computation, we should not expect medium-independence.

We seem to find an argument along these lines in Polger and Shapiro (2023). Polger and Shapiro claim that there is, of course, a notion of *abstract* computation which involves medium independence, but that it doesn't apply to physical computing systems because they don't involve abstraction:

> "Computers as abstracta would have purely mathematical or logical properties and can have properties that no physical system may have. [...] But there is an alternative view of computation that is more often applied to cognition and brains. Physical computing systems are not computers-as-abstracta. Physical computing states, transitions, and inputs or outputs are not abstract-qua-mathematical or logical but instead are concrete physical events and changes." (Polger and Shapiro 2023)

The idea seems to be this: if medium-independence is a result of the abstract properties of computation, and if physical computation does not have the abstract properties in question, then it need not be medium-independent. But why think that there is a notion of physical computation which makes no reference to abstract computational functions, only to concrete physical events and properties? Notice that this view can be found elsewhere in the literature:

> "there are two different notions of computation that are relevant here: abstract computation, as a mathematical formalism, and concrete computation, as a kind of physical causal process carried out by, for example, digital computers. [...] cognitive science and AI are primarily interested in concrete computation." (Ritchie 2011)

> "There seem to be two mutually exclusive ways to conceive of computations. One way is to view computations as a kind of formalism only; on this view, computations are abstract relations between abstract objects. Another way is to view computations as actual (types of) processes; on this view, computations are concrete processes carried out by concrete objects" (Towl 2011)

I propose that distinction between two kinds of computation, one abstract and on concrete, is mistaken: there is no such thing as 'concrete computation' which can be defined independently of abstraction. The physical *realizers* of a computation are concrete, but what makes something a physical computation is that there is an appropriate relation between the concrete entities and some computational abstraction. As mentioned previously, different theories of implementation will offer

different views of the nature of the appropriate relation and how to understand the relata: I follow Chalmers (2012) in claiming that a concrete process only counts as a physical computation if it realizes or implements some mathematical computation.

Polger and Shapiro worry that there are only two ways of thinking about abstraction, and neither of them can give us an account of physical computation. If we understand mathematical computation in terms of abstract objects, then they think it follows that there are no physical computational systems:

> "no physical system literally has the abstract computational properties that it implements, for the abstract-as-abstracta computational properties are not themselves physical properties nor specified in physical terms" (Polger and Shapiro 2023)

Notice that on this view, nothing physical is literally triangular or spherical, or even self-identical. The alternative, they propose, is that physical computations can be only abstract in the sense of being describable at a level of abstraction from the lower-level details. Polger and Shapiro reject the idea that physical systems can be computational in this sense, because they don't think this would give us the medium-independence (or multiple realizability) of a computation itself, but only relative to how we describe it:

> "[this] relies on a notion of abstraction-as-subtraction that applies primarily to descriptions or representations and is not a feature of the things represented. The subtraction of detail does not occur at the object, but at the explanation." (Polger and Shapiro 2023)

But Polger and Shapiro seems to be assuming when we say that a physical computation implements an abstract computational description, this must entail that the description in question is the mind-dependent description given by a particular individual. This is like saying that if what makes a physical object triangular or spherical is its relation to a geometric description (rather than a platonic object), then being triangular or spherical is just a property of the way we choose to describe the object. Some abstract mathematical descriptions provide definitions of objects, not merely a human-centric ways of representing objects.

I acknowledge that that once we have identified a concrete process as the implementation of a computation, then we can give an account of that process entirely physically. If we are looking to explain how a physical computational system performs as it does, we might appeal entirely to concrete entities and processes. But these concrete properties alone will not be what make the process a computational one: they will not tell us in virtue of what the system is a *computer*. For that, we need some concept of computational *abstraction* to play a role. I propose that once we have this sort of abstraction, we have medium-independence. To sum it up in a slogan, there is no

computation without abstraction, and no computational abstraction without medium-independence.[2]

3. Empirical arguments

Empirical arguments against the medium-independence of computation attempt to demonstrate that there are actual cases of physical computation which are medium-dependent. These arguments are often framed in terms of 'neural computation': the claim is that there are computations being performed by the human brain which could *not* be performed by anything other than our actual neural substrate. (This is usually framed as a claim about nomological possibility.)

The term 'neural computations' refers to the physical processes in the nervous system which carry out (or perform or implement) computations:

> "Since cognition in biological systems is a function of the nervous system, the computations that putatively explain biological cognition are carried out by the nervous system. Following the mainstream literature, we refer to them as neural computations." (Piccinini and Bahar 2013)

Chirimuuta (2022) think that neural computations are not medium independent, because the physical processes in the nervous system which perform computations are highly dependent on the particular physical properties of these processes. Chirimuuta proposes, for example, that we should assume that the spike signalling between neurons is medium-independent, because even the electrophysiological properties of the nervous system are highly dependent on the chemical properties of the nervous system, which are themselves not medium-independent:

> "inter-neuron communication mediated by spikes is not just electrical, it is chemical as well. [...] this form of information processing cannot be medium-independent. [...] The molecules are signalling in virtue of their specific material (i.e. chemical) properties, such as binding affinities, and their modulating nanoscale structures, which determine how they operate and interact with other molecules." (Chirimuuta 2022)

Similar claims are made by Cao (2022): she denies that we could substitute non-neural properties for neural properties and keep the same electrical activity. And once we move beyond spiking activity to the chemical properties of the brain, Cao claims that

[2] Michael Weisberg nicely sums up the importance of abstraction for computation as follows: "The concept of abstraction is said to permeate the entire field of computer science. [...] The idea is that the *way* things change can be independent of what things are changing." (Weisberg 2020).

this is even more obvious that the system's activity depends on its chemical and biological properties:

> "the functional organization of the brain at the cellular level exhibits characteristics that make it unlikely to be multiply realizable at that level of description. [...] Metabolic stuff is tightly intertwined with the information-processing stuff, and in the brains of creatures like us, you cannot have one without the other." (Cao 2022)

Chirimuuta and Cao agree that the actual properties of the nervous system which supposedly implement computations are, in some important sense, dependent on the neural medium itself. I don't intend to argue with this, because I don't think it is relevant to the question of whether computation is medium-dependent. They are both making a claim about the neural properties which perform physical computations, i.e. about the physical relata of the implementation relation. This tells us nothing about whether physical computation, considered as a relation between the concrete properties of the brain and abstract mathematical computations, is medium-dependent.[3]

At this point, some theorists might argue that by denying that computation is a concrete property of physical system, I am not paying enough attention to how scientists actually use the term 'neural computation'. I agree that philosophers should not ignore scientific practise, and that our philosophy of science should be informed by empirical science. But I don't think that neuroscientists' talk of "neural computation" suggests that they have in mind a notion of computation which can be defined entirely in terms of physical properties of brains.[4] It seems more likely that they are using the term 'neural computation' to refer to the neural mechanisms which others have taken to implement certain computations. And there are, of course, interesting theoretical questions to ask about the workings of our neural processes – whether or not they implement computations. We can ask what level of grain is appropriate to describe neural activity, and whether we can abstract away from certain properties of the system. But not all abstractions or higher-level claims about physical processes are thereby computational.

4. Analogicity arguments

Analogicity arguments generally acknowledge that even if there is something inherently medium-independent about *digital* computation, the same cannot be said about *analog*

---

[3] Cao (2022) is not arguing for the medium dependence of neural computation, but rather denying the multiple realizability of neural processes. But Polger and Shapiro (2023) cite her work as a reason to deny that *computation* is medium independent.

[4] As Piccinini himself notes, "In many quarters, especially neuroscientific ones, the term "computation" is used, more or less, for whatever internal processes explain cognition" (Piccinini and Scarantino 2011).

computation. I will focus here on Maley's (2021, 2023) argument for the medium-dependence of analog computation, which starts from claims about the nature of analog *representation*, and uses these to draw conclusions about analog *computation*.

Maley's argument relies on a particular definition of analog representation which rejects the 'received view' of analog representation. According to the received view (Goodman 1968), what makes a representational vehicle analog is that it is continuous rather than discrete or differentiated. According to Maley's alternative view (see also Lewis 1971, Beck 2019, Block 2023), analog representations can be continuous *or* discrete; what makes them *analog* is the structural isomorphism between their vehicles and their contents. On this alternative view, an hourglass is an analog representation of time even though the representational vehicle can be differentiated into discrete grains of sand. What is important is that as the representing property increases or decreases, it mirrors the increase or decrease of the represented property. In Peacocke's words, "Analog representation is representation of magnitudes, by magnitudes" (Peacocke 2019, 52).

It is this alternative account of analog representation which leads Maley to claim that analog representation is not medium-independent:

> "analog representation [...] requires reference to the implementing medium (and is thus medium-dependent, contra accepted views about the necessity of medium-independence for computation)" (Maley 2021)
>
> "Characterizing analog representation *qua* representation requires reference to the physical details of the system that implements those representations in a way that digital representation does not. [...] analog representation simply cannot be separated from its physical implementation in the way that digital representation can" (Maley 2021)

I'm not interested here in adjudicating between competing accounts of analog representation. I propose that even if we (a) reject the received view of analog representation in favor of Maley's 'structural isomorphism' characterization, and (b) allow that this notion of analog representation requires reference to the implementing medium, it is difficult to establish that a process operating over analog representations would itself be both computational and medium-dependent.

Notive that for Maley, a physical computational process is analog in virtue of operating over analog representation: Maley (2023) defines the concept of analog computation in terms of analog representation. We can then ask: which properties of analog representations does the computation operate over? There seem to be two options here: either analog computations operate over the properties of analog representations which individuate them as representations, or there is another way to individuate analog representations for the purposes of analog computation.

On the first of these two options, whether a physical system is an analog computer will not be a property of the physical system itself, because being an analog representation is a relational feature of the representation medium: what makes a representation analog is the structural relation between the representational medium and the distal content it represents. As a result, whether a physical system is an analog computer will depend on features beyond the physical system: the same physical system would be an analog computer in some environments and not others, and which computation it was performing would depend on which environment it was in. This results in a counter-intuitive way of individuating computers and computational processes.

On the second of these two options, analog computation operations over only those properties of analog representations which are internal to the physical system, and not the relational features which individuate them as representations. On this option, there must be some way to individuate analog representations as computational vehicles, in terms of their specific properties to which analog computations are sensitive. But what properties are these? According to Maley, the vehicles of analog representation can be discrete or continuous. So whatever analog computation is, it is a process which operates over some property of a computational vehicle which can implemented either discretely or continuously. In other words, the vehicles of analog computation are in some important sense *independent* of the medium which implements them.

I am suggesting if we define analog computation in terms of the processing of analog representations, and we commit to the 'structural isomorphism' view of what it is to be an analog representation, then we seem to face a dilemma. On the first horn of the dilemma, we end up with a concept of analog computation on which the identity conditions of physical computations and physical computers are environment-dependent, which stretches our ordinary understanding of physical computation beyond recognition. On the second horn of the dilemma, we can retain our ordinary understanding of physical computation, but the argument for medium-dependence is now much more difficult to make: it looks like analog computation must be sensitive to some property of analog representations which can be implemented in different ways (continuously or discretely).[5]

I want to close by suggest an alternative way in which we might consider analog computation to be medium-dependent, which does not involve making any claims about analog representations.[6] If we understand medium-independence in terms of Turing machines and algorithms, then there is a sense which analog physical computers do not implement algorithms. They can be characterized as implementing computational functions in input-output terms (see Pour-El 1974), but the process by which they implement a functional has no formal characterization. If this argument

[5] This is a brief account of a dilemma which I am fleshing out in more detail in other work.
[6] Chirimuuta (2022) seems to gesture in this direction.

works, it gets us a concept of medium-dependent physical computation, but not one which will be able to play any helpful role in scientific explanation. This is because there would be no way of typing tokens of analog computational processes: there is nothing which two token analog physical computations would share in virtue of which we could classify them as the same type. To the extent to which scientific explanation relies on generalization, analog physical computations would be explanatorily inert.

5. Conclusion

I have considered three recent arguments which purport to show that some subset of physical computation is medium-dependent. I have suggested that all three arguments fail. Concreta arguments start from the mistaken assumption that there is some concept of physical computation which can be specified independently of any appeal to mathematical abstraction. Empirical arguments show, at most, that the physical implementation basis of a computation can itself be medium-dependent; this is neither surprising, nor relevant to claims to the medium-independency of physical computation. The analogicity argument that I have considered here makes neither of these mistakes – and if there is an argument to be made for the medium-dependence of some physical computation, I suspect that analogicity will figure in said argument. Maley's (2021, 2023) argument for the medium-dependence of analog computation, however, relies on defining analog computation in terms of analog representation, on a particular way of understanding analog representation. I have argued that on this view of analog representation, it is difficult to see how the processes operating over these representations could be both computational and medium dependent.

References

Beck, Jacob (2019). Perception is Analog: The Argument from Weber's Law. *Journal of Philosophy* 116 (6):319-349.

Block, Ned (2023). *The Border Between Seeing and Thinking*. New York, US: OUP USA.

Cao, Rosa (2022). Multiple realizability and the spirit of functionalism. *Synthese* 200 (6):1-31.

Chalmers, David (2012). The varieties of computation: A reply. *Journal of Cognitive Science* 2012 (3):211-248.

Chirimuuta, Mazviita (2022). The Case for Medium Dependence: Comment on Neurocognitive Mechanisms by Gualtiero Piccinini. *Journal of Consciousness Studies* 29 (7-8):185-194.

Curtis-Trudel, André (2021). Implementation as Resemblance. *Philosophy of Science* 88 (5):1021-1032.

Goodman, Nelson (1968). *Languages of Art: An Approach to a Theory of Symbols*. Indianapolis: Bobbs-Merrill.

Haugeland, John (1985). *Artificial Intelligence: The Very Idea*. Cambridge: MIT Press.

Lewis, David K. (1971). Analog and digital. *Noûs* 5 (3):321-327.

Maley, Corey J. (2021). The physicality of representation. *Synthese* 199 (5-6):14725-14750.

Maley, Corey J. (2023). Analogue Computation and Representation. *British Journal for the Philosophy of Science* 74 (3):739-769.

Miłkowski, Marcin (2016). Computation and Multiple Realizability. In Vincent Müller (ed.), *Fundamental Issues of Artificial Intelligence*. Springer. pp. 29-41.

Peacocke, Christopher (2019). *The Primacy of Metaphysics*. New York, NY: OUP USA.

Piccinini, Gualtiero & Bahar, Sonya (2013). Neural Computation and the Computational Theory of Cognition. *Cognitive Science* 37 (3):453-488.

Piccinini, Gualtiero & Scarantino, Andrea (2011). Information processing, computation, and cognition. *Journal of Biological Physics* 37 (1):1-38.

Polger, Thomas W. & Shapiro, Lawrence A. (2023). The Puzzling Resilience of Multiple Realization. *Minds and Machines* 33 (2):321-345.

Pour-El, M. B. (1974). Abstract Computability and Its Relation to the General Purpose Analog Computer (Some Connections Between Logic, Differential Equations and Analog Computers). *Transactions of the American Mathematical Society*, *199*, 1–28. https://doi.org/10.2307/1996870

Ritchie, J. Brendan (2011). Chalmers on implementation and computational sufficiency. Journal of Cognitive Science 12 (4): 407-425.

Towl, Brandon N. (2011) Home, pause, or break: a critique of Chalmers on implementation. Journal of Cognitive Science 12 (4): 427-442.

Weisberg, Michael (2020). Abstraction and Representational Capacity in Computational Structures. In Levy and Godfrey-Smith (eds) *The Scientific Imagination*. Oxford University Press.

Williams, Danielle (2023). Implementation and Interpretation: A Unified Account of Physical Computation. UC Davis. ProQuest ID: WILLIAMS_ucdavis_0029D_22222. Merritt ID: ark:/13030/m57n14c2. Retrieved from https://escholarship.org/uc/item/7d0760fm